arXiv:1704.05811v1 [cs.DS] 19 Apr 2017

Online Weighted Degree-Bounded Steiner Networks via Novel Online Mixed Packing/Covering

Sina Dehghani [*†] Soheil Ehsani [*†] MohammadTaghi Hajiaghayi [*†] Vahid Liaghat [‡]
Harald Räcke [§] Saeed Seddighin [*†]

Abstract

We design the first online algorithm with poly-logarithmic competitive ratio for the *edge-weighted degree-bounded Steiner forest* (EW-DB-SF) problem and its generalized variant. We obtain our result by demonstrating a new generic approach for solving mixed packing/covering integer programs in the online paradigm. In EW-DB-SF, we are given an edge-weighted graph with a degree bound for every vertex. Given a root vertex in advance, we receive a sequence of terminal vertices in an online manner. Upon the arrival of a terminal, we need to augment our solution subgraph to connect the new terminal to the root. The goal is to minimize the total weight of the solution while respecting the degree bounds on the vertices. In the offline setting, *edge-weighted degree-bounded Steiner tree* (EW-DB-ST) and its many variations have been extensively studied since early eighties. Unfortunately, the recent advancements in the online network design problems are inherently difficult to adapt for degree-bounded problems. In particular, it is not known whether the fractional solution obtained by standard primal-dual techniques for mixed packing/covering LPs can be rounded online. In contrast, in this paper we obtain our result by using structural properties of the optimal solution, and reducing the EW-DB-SF problem to an exponential-size mixed packing/covering integer program in which every variable appears only once in covering constraints. We then design a generic *integral* algorithm for solving this restricted family of IPs.

As mentioned above, we demonstrate a new technique for solving mixed packing/covering integer programs. Define the *covering frequency* k of a program as the maximum number of covering constraints in which a variable can participate. Let m denote the number of packing constraints. We design an online *deterministic integral algorithm* with competitive ratio of $O(k \log m)$ for the mixed packing/covering integer programs. We prove the tightness of our result by providing a matching lower bound for any randomized algorithm. We note that our solution solely depends on m and k. Indeed, there can be exponentially many variables. Furthermore, our algorithm directly provides an integral solution, even if the integrality gap of the program is unbounded. We believe this technique can be used as an interesting alternative for the standard primal-dual techniques in solving online problems.

[*]University of Maryland. email: {dehghani,ehsani,seddighin}@umd.edu hajiagha@cs.umd.edu

[†]Supported in part by NSF CAREER award 1053605, NSF grant CCF-1161626, ONR YIP award N000141110662, DARPA/AFOSR grant FA9550-12-1-0423, and a Google faculty research award.

[‡]Stanford University. email: vliaghat@stanford.edu

[§]Technische Universität München. email: raecke@in.tum.de

1 Introduction

Degree-bounded network design problems comprise an important family of network design problems since the eighties. Aside from various real-world applications such as vehicle routing and communication networks [6, 33, 39], the family of degree-bounded problems has been a testbed for developing new ideas and techniques. The problem of *degree-bounded spanning tree*, introduced in Garey and Johnson's *Black Book* of NP-Completeness [30], was first investigated in the pioneering work of Fürer and Raghavachari [16] (Allerton'90). In this problem, we are required to find a spanning tree of a given graph with the goal of minimizing the maximum degree of the vertices in the tree. Let b^* denote the maximum degree in the optimal spanning tree. Fürer and Raghavachari give a parallel approximation algorithm which produces a spanning tree of degree at most $O(\log(n)b^*)$. This result was later generalized by Agrawal, Klein, and Ravi [1] to the case of degree-bounded Steiner tree (DB-ST) and degree bounded Steiner forest (DB-SF) problem. In DB-ST, given a set of terminal vertices, we need to find a subgraph of minimum maximum degree that connects the terminals. In the more generalized DB-SF problem, we are given pairs of terminals and the output subgraph should contain a path connecting each pair. Fürer and Raghavachari [17](SODA'92, J. of Algorithms'94) significantly improved the result for DB-SF by presenting an algorithm which produces a Steiner forest with maximum degree at most $b^* + 1$.

The study of DB-ST and DB-SF was the starting point of a very popular line of work on various degree-bounded network design problems; e.g. [29, 32, 28, 23, 13] and more recently [15, 14]. One particular variant that has been extensively studied was initiated by Marathe *et al.* [29] (J. of Algorithms'98): In the *edge-weighted degree-bounded spanning tree* problem, given a weight function over the edges and a degree bound b, the goal is to find a minimum-weight spanning tree with maximum degree at most b. The initial results for the problem generated much interest in obtaining approximation algorithms for the edge-weighted degree-bounded spanning tree problem [11, 10, 18, 24, 25, 26, 27, 35, 36, 37]. The groundbreaking results obtained by Goemans [19] (FOCS'06) and Singh and Lau [38] (STOC'07) settle the problem by giving an algorithm that computes a minimum-weight spanning tree with degree at most $b + 1$. Singh and Lau [28] (STOC'08) generalize their result for the *edge-weighted Steiner tree* (EW-DB-ST) and *edge-weighted Steiner forest* (EW-DB-SF) variants. They design an algorithm that finds a Steiner forest with cost at most twice the cost of the optimal solution while violating the degree constraints by at most three.

Despite these achievements in the offline setting, it was not known whether degree-bounded problems are tractable in the *online setting*. The online counterparts of the aforementioned Steiner problems can be defined as follows. The underlying graph and degree bounds are known in advance. The demands arrive one by one in an online manner. At the arrival of a demand, we need to augment the solution subgraph such that the new demand is satisfied. The goal is to be competitive against an offline optimum that knows the demands in advance.

Recently, Dehghani et al. [12] (SODA'16) explore the tractability of the Online DB-SF problem by showing that a natural greedy algorithm produces a solution in which the degree bounds are violated by at most a factor of $O(\log n)$, which is asymptotically *tight*. They analyze their algorithm using a dual fitting approach based on the combinatorial structures of the graph such as the toughness[1] factor. Unfortunately, greedy methods are not competitive for the edge-weighted variant of the problem. Hence, it seems unlikely that the approach of [12] can be generalized to EW-DB-SF.

The *online edge-weighted Steiner connectivity* problems (with no bound on the degrees) have been extensively studied in the last decades. Imase and Waxman [22] (SIAM J. D. M.'91) use a dual-fitting argument to show that the greedy algorithm has a competitive ratio of $O(\log n)$, which is also asymptotically tight.

[1]The toughness of a graph is defined as $\min_{X \subseteq V} \frac{|X|}{|\mathrm{CC}(G \setminus X)|}$; where for a graph H, $\mathrm{CC}(H)$ denotes the collection of connected components of H.

Later the result was generalized to the EW SF variant by Awerbuch *et al.* [4] (SODA'96) and Berman and Coulston [7] (STOC'97). In the past few years, various primal-dual techniques have been developed to solve the more general node-weighted variants [2, 31, 21] (SIAM'09, FOCS'11, FOCS'13), prize-collecting variants [34, 20] (ICALP'11,ICALP'14), and multicommodity buy-at-bulk [9] (FOCS'15). These results are obtained by developing various primal-dual techniques [2, 21] while generalizing the application of combinatorial properties to the online setting [31, 20, 9]. In this paper however, we develop a primal approach for solving *bounded-frequency mixed packing/covering integer programs*. We believe this framework would be proven useful in attacking other online packing and covering problems.

1.1 Our Results and Techniques

In this paper, we consider the online Steiner tree and Steiner forest problems at the presence of both edge weights and degree bounds. In the Online EW-DB-SF problem, we are given a graph $G = (V, E)$ with n vertices, edge-weight function w, degree bound b_v for every $v \in V$, and an online sequence of connectivity demands (s_i, t_i). Let w_{opt} denote the minimum weight subgraph which satisfies the degree bounds and connects all demands. Let $\rho = \frac{\max_e w(e)}{\min_{e:w(e)>0} w(e)}$[2].

Theorem 1.1 *There exists an online deterministic algorithm which finds a subgraph with total weight at most $O(\log^2 n)w_{opt}$ while the degree bound of a vertex is violated by at most a factor of $O(\log^2(n)\log(n\rho))$.*

If one favors the degree bounds over total weight, one can find a subgraph with degree-bound violation $O(\log^2(n)\frac{\log(n\rho))}{\log\log(n\rho)})$ and total cost $O(\log^2(n)\frac{\log(n\rho))}{\log\log(n\rho)})w_{opt}$.

Our technical contribution for solving the EW-DB-SF problem is twofold. First by exploiting a structural result and massaging the optimal solution, we show a formulation of the problem that falls in the restricted family of *bounded-frequency mixed packing/cover IPs*, while losing only logarithmic factors in the competitive ratio. We then design a generic online algorithm with a logarithmic competitive ratio that can solve any instance of the bounded-frequency packing/covering IPs. In what follows, we describe these contributions in detail.

Massaging the optimal solution Initiated by work of Alon *et al.* [2] on online set cover, Buchbinder and Naor developed a strong framework for solving packing/covering LPs *fractionally* online. For the applications of their general framework in solving numerous online problems, we refer the reader to the survey in [8]. Azar *et al.* [5] generalize this method for the fractional *mixed* packing and covering LPs. The natural linear program relaxation for EW-DB-SF, commonly used in the literature, is a special case of mixed packing/covering LPs: one needs to select an edge from every cut that separates the endpoints of a demand (covering constraints), while for a vertex we cannot choose more than a specific number of its adjacent edges (packing constraints). Indeed, one can use the result of Azar *et al.* [5] to find an online *fractional* solution with polylogarithmic competitive ratio. However, doing the rounding in an online manner seems very hard.

Offline techniques for solving degree-bounded problems often fall in the category of iterative and dependent rounding methods. Unfortunately, these methods are inherently difficult to adapt for an online settings since the underlying fractional solution may change dramatically in between the rounding steps. Indeed, this might be the very reason that despite many advances in the online network design paradigm in the past two decades, the natural family of degree-bounded problems has remained widely open. In this paper,

[2] Our competitive ratios have a logarithmic dependency on ρ, i.e., the ratio between largest and smallest weight. It follows from the result of [12] that one cannot obtain polylogarithmic guarantees if this ratio is not polynomially bounded

we circumvent this by reducing EW-DB-ST to a novel formulation beyond the scope of standard online packing/covering techniques and solving it using a new online integral approach.

The crux of our IP formulation is the following structural property: Let (s_i, t_i) denote the i^{th} demand. We need to augment the solution Q_{i-1} of previous steps by buying a subgraph that makes s_i and t_i connected. Let G_i denote the graph obtained by contracting the pairs of vertices s_j and t_j for every $j < i$. Note that any $(s_i - t_i)$-path in G_i corresponds to a feasible augmentation for Q_{i-1}. Some edges in G_i might be already in Q_{i-1} and therefore by using them again we can save both on the total weight and the vertex degrees. However, in Section 2 we prove that there always exists a path in G_i such that even without sharing on any of the edges in G_i and therefore paying completely for the increase in the weight and degrees, we can approximate the optimal solution up to a logarithmic factor. This in fact, enables us to have a formulation in which the covering constraints for different demands are *disentangled*. Indeed, we only have one covering constraint for each demand. Unfortunately, this implies that we have exponentially many variables, one for each possible path in G_i. This may look hopeless since the competitive factors obtained by standard fractional packing/covering methods introduced by Buchbinder and Naor [8] and Azar *et al.* [5], depend on the logarithm of the number of variables. Therefore we come up with a new approach for solving this class of mixed packing/covering integer programs (IP).

Bounded-frequency mixed packing/covering IPs We derive our result for EW-DB-ST by demonstrating a new technique for solving mixed packing/covering *integer* programs. We believe this approach could be applicable to a broader range of online problems. The integer program $\mathbb{IP}1$ describes a general mixed packing/covering IP with the set of integer variables $\mathbf{x} \in \mathbb{Z}^n_{\geq 0}$ and α. The packing constraints are described by a $m \times n$ non-negative matrix P. Similarly, the $q \times n$ matrix C describes the covering constraints. The *covering frequency* of a variable x_i is defined as the number of covering constraints in which x_i has a positive coefficient. The covering frequency of a mixed packing/covering program is defined as the maximum covering frequency of its variables.

$$\begin{aligned} \text{minimize} \quad & \alpha \ , \\ s.t. \quad & P\mathbf{x} \leq \alpha \ . \\ & C\mathbf{x} \geq 1 \ . \\ & \mathbf{x} \in \mathbb{Z}_{\geq 0}, \alpha \in \mathbb{R}_{>0} \ . \end{aligned} \tag{$\mathbb{IP}1$}$$

In the online variant of mixed packing and covering IP, we are given the packing constraints in advance. However the covering constraints arrive in an online manner. At the arrival of each covering constraint, we should *increase* the solution $\mathbf{x}$ such that it satisfies the new covering constraint. We provide a novel algorithm for solving online mixed packing/covering IPs.

Theorem 1.2 *Given an instance of the online mixed packing/covering IP, there exists a deterministic integral algorithm with competitive $O(k \log m)$, where m is the number of packing constraints and k is the covering frequency of the IP.*

We note that the competitive ratio of our algorithm is independent of the number of variables or the number of covering constraints. Indeed, there can be exponentially many variables.

Our result can be thought of as a generalization of the work of Aspnes et al. [3] (JACM'97) on virtual circuit routing. Although not explicit, their result can be massaged to solve mixed packing/covering IPs in which all the coefficients are zero or one, and the covering frequency is one. They show that such IPs

admit a $O(\log(m))$-competitive algorithms. Theorem 1.2 generalizes their result to the case with arbitrary non-negative coefficients and any bounded covering frequency.

We complement our result by proving a matching lower bound for the competitive ratio of any *randomized* algorithm. This lower bound holds even if the algorithm is allowed to return fractional solutions.

Theorem 1.3 *Any randomized online algorithm A for integral mixed packing and covering is $\Omega(k \log m)$-competitive, where m denotes the number of packing constraints, and k denotes the covering frequency of the IP. This even holds if A is allowed to return a fractional solution.*

As mentioned before, Azar *et al.* [5] provide a fractional algorithm for mixed packing/covering LPs with competitive ratio of $O(\log m \log d)$ where d is the maximum number of variables in a single constraint. They show an almost matching lower bound for deterministic algorithms. We distinguish two advantages of our approach compared to that of Azar *et al*:

- The algorithm in [5] outputs a *fractional* competitive solution which then needs to be rounded online. For various problems such as Steiner connectivity problems, rounding a solution online is very challenging, even if offline rounding techniques are known. Moreover, the situation becomes hopeless if the integrality gap is unbounded. However, for bounded-frequency IPs, our algorithm directly produces an integral competitive solution even if the integrality gap is large.
- Azar *et al.* find the best competitive ratio with respect to the number of packing constraints and the size of constraints. Although these parameters are shown to be bounded in several problems, in many problems such as connectivity problems and flow problems, formulations with exponentially many variables are very natural. Our techniques provide an alternative solution with a tight competitive ratio, for formulations with bounded covering frequency.

1.2 Preliminaries

Let $G = (V, E)$ be an undirected graph of size n ($|V| = n$). Let $w : E \to \mathbb{Z}_{>0}$ be a function denoting the edge weights. For a subgraph $H \subseteq G$, we define $w(H) := \sum_{e \in E(H)} w(e)$. For every vertex $v \in V$, let $b_v \in \mathbb{Z}_{>0}$ denote the degree bound of v. Let $\deg_H(v)$ denote the degree of vertex v in subgraph H. We define the load $l_H(v)$ of vertex v w.r.t. H as $\deg_H(v)/b_v$. In DB-SF we are given graph G, degree bounds, and k connectivity demands. Let σ_i denote the i-th demand. The i-th demand is a pair of vertices $\sigma_i = (s_i, t_i)$, where $s_i, t_i \in V$. In DB-SF the goal is to find a subgraph $H \subseteq G$ such that for each demand σ_i, s_i is connected to t_i in H, for every vertex $v \in V$, $l_H(v) \leq 1$, and $w(H)$ is minimized. In this paper without loss of generality we assume the demand endpoints are distinct vertices with degree one in G and degree bound infinity.

In the online variant of the problem, we are given graph G and degree bounds in advance. However the sequence of demands are given one by one. At arrival of demand σ_i, we are asked to provide a subgraph H_i, such that $H_{i-1} \subseteq H_i$ and s_i is connected to t_i in H_i.

The following integer program is a natural mixed packing and covering integer program for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST. Let $\mathcal{S}$ denote the collection of subsets of vertices that separate the endpoints of at least one demand. For a set of vertices S, let $\delta(S)$ denote the set of edges with exactly one endpoint in S. In $\mathbb{SF}_\mathbb{IP}$, for an edge e, $x_e = 1$ indicates that we include e in the solution while $x_e = 0$ indicates otherwise. The variable α indicates an upper bound on the violation of the load of every vertex and an upper bound on the violation of the weight. The first set of constraints ensures that the load of a vertex is upper bounded by α. The second constraint ensures that the violation for

the weight is upper bounded by α. The third set of constraints ensures that the endpoints of every demand are connected.

$$\begin{aligned}
&\text{minimize} \quad \alpha \ . && (\mathbb{SF_IP})\\
&\forall v \in V \quad \frac{1}{b_v} \sum_{e \in \delta(\{v\})} x_e \leq \alpha \ . && (1)\\
&\qquad\qquad \frac{1}{w_{\text{opt}}} \sum_{e \in E} w(e) x_e \leq \alpha \ . && (2)\\
&\forall S \subseteq \mathcal{S} \quad \sum_{e \in \delta(S)} x_e \geq 1 \ . && (3)\\
&\qquad\qquad x_e \in \{0,1\}, \alpha \in \mathbb{Z}_{>0} \ .
\end{aligned}$$

1.3 Overview of the Paper

We begin Section 2 by providing an online bounded frequency mixed packing/covering IP for EW-DB-SF. Further we prove that this formulation has plausible structures. In Section 3 we provide an online deterministic algorithm for online bounded frequency mixed packing/covering IPs. In Section 4 we merge the IP formulation in Section 2 and the techniques in Section 3 to obtain online polylogarithmic-competitive algorithms for EW-DB-SF. Finally in Section A we complement our algorithm for online bounded frequency mixed packing/covering IPs by providing a matching lower bound for the competitive ratio of any randomized algorithm.

2 Finding the Right Integer Program

In this section we design an online mixed packing and covering integer program for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST. We show this formulation is near optimal, i.e. any $f-$approximation for this formulation, implies an $O(f \log^2 n)$-approximation for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST. In Section 4 we show there exists an online algorithm that finds an $O(\log^3 n)$-approximation of w_{opt} and violates degree bounds by $O(\log^3 n \log w_{\text{opt}})$, where w_{opt} denotes the optimal weight.

First we define some notations. For a sequence of demands $\sigma = \langle (s_1, t_1), \ldots, (s_k, t_k) \rangle$, we define $R_\sigma(i)$ to be a set of i edges, connecting the endpoints of the first i demands. In particular $R_\sigma(i) := \bigcup_{j=1}^{i} e(s_j, t_j)$, where $e(s_j, t_j)$ denotes a direct edge from s_j to t_j. Moreover, we say subgraph H_i satisfies the connectivity of demand $\sigma_i = (s_i, t_i)$, if s_i and t_i are connected in graph $H_i \cup R_\sigma(i-1)$. Let $\mathcal{H}_i$ denote the set of all subgraphs that satisfy the connectivity of demand σ_i. In $\mathbb{PC_IP}$ variable α denotes the violation in the packing constraints. Furthermore for every subgraph $H \subseteq G$ and demand σ_i, there exists a variable $x_H^i \in \{0,1\}$. $x_H^i = 1$ indicates we add the edges of H to the existing solution, at arrival of demand σ_i. The first set of constraints ensure the degree-bounds are not violated more than α. The second constraint ensures the weight is not violated by more than α. The third set of constraints ensure the endpoints of every demand are connected.

$$\text{minimize} \quad \alpha \ . \tag{$\mathbb{PC}_\mathbb{IP}$}$$

$$\forall v \in V \qquad \frac{1}{b_v} \sum_{i=1}^{k} \sum_{H \subseteq G} \deg_H(v) x_H^i \leq \alpha \ . \tag{4}$$

$$\frac{1}{w_{\text{opt}}} \sum_{i=1}^{k} \sum_{H \subseteq G} w(H) x_H^i \leq \alpha \ . \tag{5}$$

$$\forall \sigma_i \qquad \sum_{H \in \mathcal{H}_i} x_H^i \geq 1 \ . \tag{6}$$

$$\forall H \subseteq G, 1 \leq i \leq k \qquad \mathbf{x}_H^i \in \{0,1\} \ .$$

$$\alpha > 0 \ .$$

We are considering the online variant of the mixed packing and covering program. We are given the packing Constraints (4) and (5) in advance. At arrival of demand σ_i, the corresponding covering Constraint (6) is added to the program. We are looking for an online solution which is feasible at every online stage. Moreover the variables x_H should be monotonic, i.e. once an algorithm sets $x_H = 1$ for some H, the value of x_H is 1 during the rest of the algorithm. Figure (1) illustrates an example which indicates the difference between the solutions of $\mathbb{PC}_\mathbb{IP}$ and $\mathbb{SF}_\mathbb{IP}$.

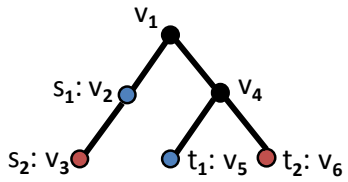

Figure 1: An example where every vertex has degree-bound 3 and every edge has weight 1. The first demand is (v_2, v_5) and the second demand is (v_3, v_6). The optimal solution for $\mathbb{SF}_\mathbb{IP}$ is a subgraph, say H, with the set of all edges and vertices, i.e. $H = G$. However an optimal solution for $\mathbb{PC}_\mathbb{IP}$ is: Two subgraphs H_1 for the first request which has edges $\{e(v_1, v_2), e(v_1, v_4), e(v_4, v_5)\}$ and H_2 for the second request which has edges $\{e(v_2, v_3), e(v_4, v_5), e(v_4, v_6)\}$. Note that $w(H) = 5$ and $w(H_1) + w(H_2) = 6$, since we have edge $e(v_4, v_5)$ in both H_1 and H_2. Moreover the number of edges incident with v_4 in the solution of $\mathbb{PC}_\mathbb{IP}$ is 4, i.e. $\deg_{H_1}(v_4) + \deg_{H_2}(v_4) = 4$.

Let opt_{IP} and opt_{SF} denote the optimal solutions for $\mathbb{PC}_\mathbb{IP}$ and $\mathbb{SF}_\mathbb{IP}$, respectively. Lemma 2.1 shows given an online solution for $\mathbb{PC}_\mathbb{IP}$ we can provide a feasible online solution for $\mathbb{SF}_\mathbb{IP}$ of cost opt_{IP}.

Lemma 2.1 *Given a feasible solution $\{\mathbf{x}, \alpha\}$ for $\mathbb{PC}_\mathbb{IP}$, there exists a feasible solution $\{\mathbf{x}', \alpha\}$ for $\mathbb{SF}_\mathbb{IP}$.*

In the rest of this section, we show that we do not lose much by changing $\mathbb{SF}_\mathbb{IP}$ to $\mathbb{PC}_\mathbb{IP}$.
In particular Lemma G.1 shows $\text{opt}_{\text{IP}} \leq O(\log^2 n)\text{opt}_{\text{SF}}$. To this end, we first define the *connective* list of subgraphs for a graph G, a forest F, and a list of demands σ. We then prove an existential lemma for such a list of subgraphs with a desirable property for any $\langle G, F, \sigma \rangle$. With that in hand, we prove $\text{opt}_{\text{IP}} \leq O(\log n)\text{opt}_{\text{SF}}$ in Lemma G.1.

Given a graph G, a list of demands $\sigma = \langle (s_1, t_1), (s_2, t_2) \ldots, (s_k, t_k) \rangle$, and a forest of G, F we define a *connective* list of subgraphs for $\langle G, F, \sigma \rangle$ in the following way:

Definition 2.2 *Let $Q = \langle Q_1, Q_2, Q_3, \ldots, Q_k \rangle$ be a list of k subgraphs of F. We say Q is a connective list of subgraphs for $\langle G, F, \sigma \rangle$ iff for every $1 \leq i \leq k$ there exists no cut disjoint from Q_i that separates s_i from t_i, but does not separate any s_j from t_j for $j < i$.*

The intuition behind the definition of connective subgraphs is the following: If Q is a connective list of subgraphs for an instance $\langle G, F, \sigma \rangle$ then for every i we are guaranteed that the union of all subgraphs $\cup_{j=1}^{i} Q_i$ connects s_i to t_i. In Lemma 2.3 we show for every $\langle G, F, \sigma \rangle$, there exists a connective list of subgraphs for $\langle G, F, \sigma \rangle$, such that each edge of F appears in at most $O(\log n)$ subgraphs of Q.

Lemma 2.3 *Let G be a graph and F be a forest in G. If σ is a collection of k demands $\langle (s_1, t_1), (s_2, t_2) \ldots, (s_k, t_k) \rangle$, then there exists a connective list of subgraphs $Q = \langle Q_1, Q_2, \ldots, Q_k \rangle$ for $\langle G, F, \sigma \rangle$ such that every edge of F appears in at most $3 \log |V(F)|$ Q_i's.*

Finally, we leverage Lemma 2.3 to show $\text{opt}_{\text{IP}} \leq O(\log n)\text{opt}_{\text{SF}}$.

Lemma 2.4 *$opt_{IP} \leq O(\log n) opt_{SF}$.*

Finally, we leverage Lemma 2.3 to show $\text{opt}_{\text{IP}} \leq O(\log^2 n)\text{opt}_{\text{SF}}$.

Lemma 2.5 *$opt_{IP} \leq O(\log^2 n) opt_{SF}$.*

This shows we can use $\mathbb{PC}_\mathbb{IP}$ as an online mixed packing/covering IP to obtain an online solution for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST losing a factor of $O(\log^2 n)$.
In Section 4 we show this formulation is an online bounded frequency mixed packing/covering IP, thus we leverage our technique for such IPs to obtain a polylogarithmic-competitive algorithm for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST.

3 Online Bounded Frequency Mixed Packing/Covering IPs

In this section we consider bounded frequency online mixed packing and covering integer programs. For every online mixed packing and covering IP with covering frequency k, we provide an online algorithm that violates each packing constraint by at most a factor of $O(k \log m)$, where m is the number of packing constraints. We note that this bound is independent of the number of variables, the number of covering constraints, and the coefficients of the mixed packing and covering program. Moreover the algorithm is for integer programs, which implies obtaining an integer solution does not rely on (online) rounding.

In particular we prove there exists an online $O(k \log m)$-competitive algorithm for any mixed packing and covering IP such that every variable has covering frequency at most k, where the covering frequency of a variable x_r is the number of covering constraints with a non-zero coefficient for x_r.

We assume that all variables are binary. One can see this is without loss of generality as long as we know every variable $x_r \in \{1, 2, 3, \ldots, 2^l\}$. Since we can replace x_r by l variables $y_r^1, \ldots, y_r^l$ denoting the digits of x_r and adjust coefficients accordingly. Furthermore, for now we assume that the optimal solution for the given mixed packing and covering program is 1. In Theorem 3.3 we prove that we can use a doubling technique to provide an $O(k \log m)$-competitive solution for online bounded frequency mixed packing and covering programs with any optimal solution. The algorithm is as follows. We maintain a family of subsets $\mathcal{S}$. Initially $\mathcal{S} = \emptyset$. Let $\mathcal{S}(j)$ denote $\mathcal{S}$ at arrival of C_{j+1}. For each covering constraint C_{j+1}, we find a subset of variables S_{j+1} and add S_{j+1} to $\mathcal{S}$. We find S_{j+1} in the following way. For each set of variables S, we define a cost function $\tau_S(\mathcal{S}(j))$ according to our current $\mathcal{S}$ at arrival of C_{j+1}. We find a set S_{j+1} that satisfies C_{j+1} and minimizes $\tau_S(\mathcal{S}(j))$. More precisely we say a set of variables S satisfies C_{j+1} if

- $\sum_{x_r \in S} C_{j+1,r} x_r \geq 1$, where $C_{j+1,r}$ denotes the coefficient of C_{j+1} for x_r.
- For each packing constraint P_i, $\sum_{x_r \in S} \frac{1}{k} P_{ir} \leq 1$.

Now we add S_{j+1} to $\mathcal{S}$ and for every $x_r \in S_{j+1}$, we set $x_r = 1$. We note that there always exists a set S that satisfies C_{j+1}, since we assume there exists an optimal solution with value 1. Setting S to be the set of all variables with value one in an optimal solution which have non-zero coefficient in C_{j+1}, satisfies C_{j+1}. It only remains to define $\tau_S(\mathcal{S}(j))$. But before that we need to define $\Delta_i(S)$ and $F_i(\mathcal{S}(j))$. For packing constraint P_i and subset of variables S, we define $\Delta_i(S)$ as $\Delta_i(S) := \sum_{x_r \in S} \frac{1}{k} P_{ir}$. For packing constraint P_i and $\mathcal{S}(j)$, let

$$F_i(\mathcal{S}(j)) := \sum_{S \in \mathcal{S}(j)} \Delta_i(S) \ . \tag{7}$$

Now let $\tau_S(\mathcal{S}(j)) = \sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j)) + \Delta_i(S)} - \rho^{F_i(\mathcal{S}(j))}$, where $\rho > 1$ is a constant to be defined later.

Algorithm 1

Input: Packing constraints P, and an online stream of covering constraints $C_1, C_2, \ldots$.
Output: A feasible solution for online bounded frequency mixed packing/covering.
Offline Process:

1: Initialize $\mathcal{S} \leftarrow \emptyset$.

Online Scheme; assuming a covering constraint C_{j+1} is arrived:

1: $S_{j+1} \leftarrow \arg\min_S \{\tau_S(\mathcal{S}(j)) \mid S \text{ satisfies } C_{j+1}\}$.
2: **for all** $x_r \in S_{j+1}$ **do**
3: $\quad x_r \leftarrow 1$.

Let $\mathbf{x}^*$ be an optimal solution, and $\mathbf{x}^*(j)$ denote its values at online stage j. We define $G_i(j)$ as

$$G_i(j) := \sum_{l=1}^{j} \sum_{r: C_{lr} > 0} \frac{1}{k} x_r^* P_{ir} \ . \tag{8}$$

Now we define a potential function Φ_j for online stage j.

$$\Phi_j = \sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j))} (\gamma - G_i(j)) \ , \tag{9}$$

where $\rho, \gamma > 1$ are constants to be defined later.

Lemma 3.1 *There exist constants ρ and γ, such that Φ_j is non-increasing.*

Proof. We find ρ and γ such that $\Phi_{j+1} - \Phi_j \leq 0$. By the definition of Φ_j,

$$\Phi_{j+1} - \Phi_j = \sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j+1))} (\gamma - G_i(j+1)) - \rho^{F_i(\mathcal{S}(j))} (\gamma - G_i(j)) \ . \tag{10}$$

By Equation (7), $\rho^{F_i(\mathcal{S}(j+1))} - \rho^{F_i(\mathcal{S}(j))} = \rho^{F_i(\mathcal{S}(j)) + \Delta_i(S)} - \rho^{F_i(\mathcal{S}(j))}$. Moreover by Equation (8), $(\gamma - G_i(j+1)) - (\gamma - G_i(j)) = -\sum_{r: C_{j+1,r} > 0} \frac{1}{k} x_r^* P_{ir}$. For simplicity of notation we define $B_i(j+1) := \sum_{r: C_{j+1,r} > 0} \frac{1}{k} x_r^* P_{ir}$. Thus we can write Equation (10) as:

$$\Phi_{j+1} - \Phi_j = \sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j+1))} (\gamma - G_i(j) - B_i(j+1)) - \rho^{F_i(\mathcal{S}(j))} (\gamma - G_i(j)) \tag{11}$$

$$
\begin{aligned}
&= \sum_{i=1}^{m} (\gamma - G_i(j))(\rho^{F_i(\mathcal{S}(j))+\Delta_i(S)} - \rho^{F_i(\mathcal{S}(j))}) - \rho^{F_i(\mathcal{S}(j+1))} B_i(j+1) && \text{Since } G_i(j) \geq 0 \\
&\leq \sum_{i=1}^{m} \gamma(\rho^{F_i(\mathcal{S}(j))+\Delta_i(S)} - \rho^{F_i(\mathcal{S}(j))}) - \rho^{F_i(\mathcal{S}(j+1))} B_i(j+1) && F_i(\mathcal{S}(j+1)) \geq F_i(\mathcal{S}(j)) \\
&\leq \sum_{i=1}^{m} \gamma(\rho^{F_i(\mathcal{S}(j))+\Delta_i(S)} - \rho^{F_i(\mathcal{S}(j))}) - \rho^{F_i(\mathcal{S}(j))} B_i(j+1) \ .
\end{aligned}
$$

Now according to the algorithm for each subset of variables S' such that $\sum_{x_r \in S'} C_{j+1}(x_r) \geq 1$, either $\tau_S(\mathcal{S}(j)) \leq \tau_{S'}(\mathcal{S}(j))$ or there exists a packing constraint P_i such that $\Delta_i(S') > 1$. In $B_i(j+1)$, we are considering variables x_r such that $x_e^* = 1$, thus for every P_i, $B_i(j+1) \leq 1$. Therefore setting S' to be the set of variables x_r such that $x_r^* = 1$ and $C_{j+1,r} > 0$, we have $\tau_S(\mathcal{S}(j)) \leq \tau_{S'}(\mathcal{S}(j))$. Thus $\sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j))+\Delta_i(S)} - \rho^{F_i(\mathcal{S}(j))} \leq \sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j))+B_i(j+1)} - \rho^{F_i(\mathcal{S}(j))}$. Therefore we can rewrite Inequality (11) as

$$
\begin{aligned}
\Phi_{j+1} - \Phi_j &\leq \sum_{i=1}^{m} \gamma(\rho^{F_i(\mathcal{S}(j))+B_i(j+1)} - \rho^{F_i(\mathcal{S}(j))}) - \rho^{F_i(\mathcal{S}(j))} B_i(j+1) && (12) \\
&= \sum_{i=1}^{m} \rho^{F_i(\mathcal{S}(j))}(\gamma \rho^{B_i(j+1)} - \gamma - B_i(j+1)) \ .
\end{aligned}
$$

We would like to find ρ and γ such that Φ_j is non-increasing. We find ρ and γ such that for each packing constraint P_i, $\gamma \rho^{B_i(j+1)} - \gamma - B_i(j+1) \leq 0$. Thus

$$
\begin{aligned}
\gamma \rho^{B_i(j+1)} - \gamma &\leq B_i(j+1) && \text{Since } 0 \leq B_i(j+1) \leq 1 && (13) \\
\gamma \rho B_i(j+1) - \gamma &\leq B_i(j+1) && \text{By simplifying} && (14) \\
\rho &\leq 1 + 1/\gamma \ . && && (15)
\end{aligned}
$$

Thus if we set $\rho \leq 1 + 1/\gamma$, Φ_j is non-increasing, as desired. □

Now we prove Algorithm 1 obtains a solution of at most $O(k \log m)$.

Lemma 3.2 *Given an online bounded frequency mixed packing covering IP with optimal value 1, there exists a deterministic integral algorithm with competitive ratio $O(k \log m)$, where m is the number of packing constraints and k is the covering frequency of the IP.*

Proof. By Lemma 3.1 for each stage j, $\Phi_{j+1} \leq \Phi_j$. Therefore $\Phi_j \leq \Phi_0 = \gamma m$. Thus for each packing constraint P_i,

$$
\rho^{F_i(\mathcal{S}(j))}(\gamma - G_i(j)) \leq \gamma m \ . \tag{16}
$$

Thus,

$$
\rho^{F_i(\mathcal{S}(j))} \leq \frac{\gamma m}{(\gamma - G_i(j))} \leq \frac{\gamma m}{\gamma - 1} \ . \qquad \text{Since } G_i(j) \leq 1 \tag{17}
$$

Thus we can conclude

$$
F_i(\mathcal{S}(j)) \in O(\log m) \ . \tag{18}
$$

By definition of $F_i(\mathcal{S}(j))$, $F_i(\mathcal{S}(j)) = \sum_{S \in \mathcal{S}(j)} \Delta_i(S) = \sum_{S \in \mathcal{S}(j)} \sum_{x_r \in S} \frac{1}{k} P_{ir}$. Since each variable x_r is present in at most k sets, $\frac{1}{k} P_i \cdot \mathbf{x}(j) \leq F_i(\mathcal{S}(j))$. Thus by Inequality (18) $P_i \mathbf{x}(j) \in O(k \log m)$, which completes the proof. □

In Theorem 3.3 we prove there exists an online $O(k \log m)$-competitive algorithm for bounded frequency online mixed packing and covering integer programs with any optimal value.

Theorem 3.3 *Given an instance of the online mixed packing/covering IP, there exists a deterministic integral algorithm with competitive ratio $O(k \log m)$, where m is the number of packing constraints and k is the covering frequency of the IP.*

4 Putting Everything Together

In this section we consider the online mixed packing/covering formulation discussed in Section 2 for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST $\mathbb{PC}_\mathbb{IP}$. In this section we show this formulation is an online bounded frequency mixed packing/covering IP. Therefore we our techniques discussed in Section 3 to obtain a polylogarithmic-competitive algorithm for ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST.

First we assume we are given the optimal weight w_{opt} as well as degree bounds. We can obtain the following theorem.

Theorem 4.1 *Given the optimal weight w_{opt}, there exists an online deterministic algorithm which finds a subgraph with total weight at most $O(\log^3 n) w_{opt}$ while the degree bound of a vertex is violated by at most a factor of $O(\log^3(n))$.*

Proof. By Lemma 2.1, given a feasible online solution for $\mathbb{PC}_\mathbb{IP}$ with violation α, we can provide an online solution for $\mathbb{SF}_\mathbb{IP}$ with violation α. Moreover by Lemma G.1, $\text{opt}_{\text{IP}} \leq O(\log^2 n)\text{opt}_{\text{SF}}$. Thus given an online solution for $\mathbb{PC}_\mathbb{IP}$ with competitive ratio f, there exists an $O(f \log n)$-competitive algorithm for ONLINE DEGREE-BOUNDED STEINER FOREST. We note that in $\mathbb{PC}_\mathbb{IP}$ we know the packing constraints in advance. In addition every variable x_H^i has non-zero coefficient only in the covering constraint corresponding to connectivity of the i-th demand endpoints, i.e. the covering frequency of every variable is 1. Therefore by Theorem 3.3 there exists an online $O(\log m)$-competitive solution for $\mathbb{PC}_\mathbb{IP}$, where m is the number of packing constraints, which is $n+1$. Thus there exists an online $O(\log^3 n)$-competitive algorithm for ONLINE DEGREE-BOUNDED STEINER FOREST. This means the violation for both degree bounds and weight is of $O(\log^3 n)$. □

Now we assume we are not given w_{opt}. The following theorems directly hold by applying the doubling technique mentioned in Section D.

Theorem 4.2 *There exists an online deterministic algorithm which finds a subgraph with total weight at most $O(\log^3 n) w_{opt}$ while the degree bound of a vertex is violated by at most a factor of $O(\log^3(n) \log(w_{opt}))$.*

Moreover if one favors the degree bound over total weight we obtain the following bounds.

Theorem 4.3 *There exists an online deterministic algorithm which finds a subgraph with total weight at most $O(\log^3 n \frac{\log w_{opt}}{\log \log w_{opt}}) w_{opt}$ while the degree bound of a vertex is violated by at most a factor of $O(\log^3 n \frac{\log w_{opt}}{\log \log w_{opt}})$.*

References

[1] A. Agrawal, P. N. Klein, and R. Ravi. How tough is the minimum-degree steiner tree?: A new approximate min-max equality. *Technical Report CS-91-49, Brown University*, 1991.

[2] N. Alon, B. Awerbuch, Y. Azar, N. Buchbinder, and J. Naor. The online set cover problem. *SIAM Journal on Computing*, 39(2):361–370, 2009.

[3] J. Aspnes, Y. Azar, A. Fiat, S. Plotkin, and O. Waarts. On-line routing of virtual circuits with applications to load balancing and machine scheduling. *Journal of the ACM (JACM)*, 44(3):486–504, 1997.

[4] B. Awerbuch, Y. Azar, and Y. Bartal. On-line generalized steiner problem. In *Proceedings of the seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 68–74, 1996.

[5] Y. Azar, U. Bhaskar, L. Fleischer, and D. Panigrahi. Online mixed packing and covering. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 85–100. SIAM, 2013.

[6] F. Bauer and A. Varma. Degree-constrained multicasting in point-to-point networks. In *INFOCOM'95. Fourteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Bringing Information to People. Proceedings. IEEE*, pages 369–376. IEEE, 1995.

[7] P. Berman and C. Coulston. On-line algorithms for steiner tree problems. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 344–353, 1997.

[8] N. Buchbinder and J. Naor. The design of competitive online algorithms via a primal: dual approach. *Foundations and Trends® in Theoretical Computer Science*, 3(2–3):93–263, 2009.

[9] D. Chakrabarty, A. Ene, R. Krishnaswamy, and D. Panigrahi. Online buy-at-bulk network design. In *FOCS*, 2015.

[10] K. Chaudhuri, S. Rao, S. Riesenfeld, and K. Talwar. A push-relabel algorithm for approximating degree bounded msts. In *Proceedings of the 33rd international conference on Automata, Languages and Programming-Volume Part I*, pages 191–201. Springer-Verlag, 2006.

[11] K. Chaudhuri, S. Rao, S. Riesenfeld, and K. Talwar. What would edmonds do? augmenting paths and witnesses for degree-bounded msts. *Algorithmica*, 55(1):157–189, 2009.

[12] S. Dehghani, S. Ehsani, M. Hajiaghayi, and V. Liaghat. Online degree-bounded steiner network design. In *SODA*, 2016.

[13] A. Ene and A. Vakilian. Improved approximation algorithms for degree-bounded network design problems with node connectivity requirements. *STOC*, 2014.

[14] A. Ene and A. Vakilian. Improved approximation algorithms for degree-bounded network design problems with node connectivity requirements. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 754–763. ACM, 2014.

[15] T. Fukunaga and R. Ravi. Iterative rounding approximation algorithms for degree-bounded node-connectivity network design. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 263–272. IEEE, 2012.

[16] M. Fürer and B. Raghavachari. An NC approximation algorithm for the minimum degree spanning tree problem. In *Allerton Conf. on Communication, Control and Computing*, pages 274–281, 1990.

[17] M. Fürer and B. Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. *Journal of Algorithms*, 17(3):409–423, 1994.

[18] M. X. Goemans. Minimum bounded degree spanning trees. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 273–282. IEEE, 2006.

[19] M. X. Goemans. Minimum bounded degree spanning trees. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 273–282, 2006.

[20] M. Hajiaghayi, V. Liaghat, and D. Panigrahi. Near-optimal online algorithms for prize-collecting steiner problems. In *Automata, Languages, and Programming*, pages 576–587. 2014.

[21] M. T. Hajiaghayi, V. Liaghat, and D. Panigrahi. Online node-weighted steiner forest and extensions via disk paintings. In *Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on*, pages 558–567, 2013.

[22] M. Imase and B. M. Waxman. Dynamic Steiner tree problem. *SIAM Journal on Discrete Mathematics*, 4(3):369–384, 1991.

[23] R. Khandekar, G. Kortsarz, and Z. Nutov. On some network design problems with degree constraints. *Journal of Computer and System Sciences*, 79(5):725–736, 2013.

[24] P. N. Klein, R. Krishnan, B. Raghavachari, and R. Ravi. Approximation algorithms for finding low-degree subgraphs. *Networks*, 44(3):203–215, 2004.

[25] J. Könemann and R. Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 537–546. ACM, 2000.

[26] J. Könemann and R. Ravi. Primal-dual meets local search: approximating mst's with nonuniform degree bounds. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 389–395. ACM, 2003.

[27] L. C. Lau, J. Naor, M. R. Salavatipour, and M. Singh. Survivable network design with degree or order constraints. *SIAM Journal on Computing*, 39(3):1062–1087, 2009.

[28] L. C. Lau and M. Singh. Additive approximation for bounded degree survivable network design. *SIAM Journal on Computing*, 42(6):2217–2242, 2013.

[29] M. V. Marathe, R. Ravi, R. Sundaram, S. Ravi, D. J. Rosenkrantz, and H. B. Hunt III. Bicriteria network design problems. *Journal of Algorithms*, 28(1):142–171, 1998.

[30] R. G. Michael and S. J. David. Computers and intractability: a guide to the theory of np-completeness. *WH Freeman & Co., San Francisco*, 1979.

[31] J. Naor, D. Panigrahi, and M. Singh. Online node-weighted steiner tree and related problems. In *Foundations of Computer Science (FOCS), 2011 IEEE 52nd Annual Symposium on*, pages 210–219, 2011.

[32] Z. Nutov. Degree-constrained node-connectivity. In *LATIN 2012: Theoretical Informatics*, pages 582–593. 2012.

[33] C. A. Oliveira and P. M. Pardalos. A survey of combinatorial optimization problems in multicast routing. *Computers & Operations Research*, 32(8):1953–1981, 2005.

[34] J. Qian and D. P. Williamson. An o (logn)-competitive algorithm for online constrained forest problems. In *Automata, Languages and Programming*, pages 37–48. 2011.

[35] B. Raghavachari. Algorithms for finding low degree structures. In *Approximation algorithms for NP-hard problems*, pages 266–295. PWS Publishing Co., 1996.

[36] R. Ravi, M. V. Marathe, S. Ravi, D. J. Rosenkrantz, and H. B. Hunt III. Approximation algorithms for degree-constrained minimum-cost network-design problems. *Algorithmica*, 31(1):58–78, 2001.

[37] R. Ravi and M. Singh. Delegate and conquer: An lp-based approximation algorithm for minimum degree msts. In *Automata, Languages and Programming*, pages 169–180. Springer, 2006.

[38] M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 661–670, 2007.

[39] S. Voß. Problems with generalized steiner problems. *Algorithmica*, 7(1):333–335, 1992.

A Lower Bound

We present a (randomized) instance I for mixed packing and covering linear programs with the following parameters. I consists of m packing constraints, $(2m-2)d$ variables and a variable is only contained in $\log_2 d$ covering constraints. There exists an optimum integral solution for I with violation $\alpha = 1$. Any (fractional) online algorithm A incurs an expected violation of at least $\Omega(\log_2 m \cdot \log_2 d)$, where the expectation is w.r.t. the randomized construction of I and w.r.t. the random choices of A in case A uses randomization. This gives Theorem 1.3.

For the following description of the instance we assume that m and d are powers of 2. Consider a binary tree T with m leaf nodes. For each edge e in this tree we introduce d variables, and we denote the set of variables for edge e with X_e. For each leaf node v we introduce the packing constraint $\sum_{e \in P_v} \sum_{x \in X_e} x \leq \alpha$, where P_v denotes the path from the root r to v in the tree.

The covering constraints are constructed in an online manner according to a random root-to-leaf path P_ℓ in the tree (ℓ is a random leaf node). For a non-leaf node v on this path (starting at the root and ending at the parent of ℓ) we construct $\log_2 d$ covering constraint only involving variables from $X_{e_L} \cup X_{e_R}$, where e_L and e_R denote the child-edges of v. This sequence of covering constraints is constructed according to the following lemma. For a set X of variables we use $w(X)$ to denote the total value assigned to these variables by the online algorithm.

Lemma A.1 *Given two sets X_L and X_R of variables, each of cardinality d. There is a randomized sequence of $\log_2 d$ covering constraints over variables from $X_L \cup X_R$ such that any online algorithm has $E[w(X_L \cup X_R)] \geq \frac{1}{2} \log_2 d$, after fulfilling these constraints.*

Furthermore, there exist variables $x_\ell \in X_L, x_r \in X_R$ such that setting either of these variables to one already fulfills all constraints.

Proof. Define $A_L \subseteq X_L$ and $A_R \subseteq X_R$, as a set of *active variables* of X_L and X_R, respectively. Initially all variables in X_L and X_R are active, i.e., $A_L = X_L$ and $A_R = X_R$.

The constraints are constructed in $\log_2 d$ rounds. In the beginning of the i-th round ($i \in \{0, \ldots, \log_2 d - 1\}$) $|A_L| = |A_R| = d/2^i$. We offer a covering constraint on the current set of active variables, i.e., we offer constraint $\sum_{x \in A_L} x + \sum_{x \in A_R} x \geq 1$. Then we remove 50% of the elements from A_L, and 50% of the elements from A_R, at random, i.e., from each set we remove a random subset of cardinality $d/2^{i+1}$.

After fulfilling the covering constraint for the i-th round, we have $w(A_L) + w(A_R) \geq 1$. Removing random subsets from A_L and A_R, removes variables of expected total weight at least $w(A_L)/2 + w(A_R)/2 \geq 1/2$. Hence, the total expected weight removed from active sets during all rounds is at least $1/2 \cdot \log_2 d$. This gives the bound on the expected weight of variables.

Note that any variable that is active in the final round is contained in every constraint; setting any of these variables to one fulfills all constraints. Since, neither A_L nor A_R is empty in the final round the lemma follows. □

Note that by this construction an optimal offline algorithm can fulfill the covering constraints of a node v by setting a single variable from set X_e to 1, where e is the child-edge of v that is *not* on the path P_ℓ. Then, along any root-to-leaf path at most one set X_e contains a variable that is set to 1, and, hence, the maximum violation is $\alpha = 1$.

For the online algorithm we show that $E[\sum_{e \in P_\ell} w(X_e)] \geq \frac{1}{4}(\log_2 m - 1) \cdot \log_2 d$, which means that in expectation the maximum violation is at least $\Omega(\log m \cdot \log d)$. To see this, we represent the path P_ℓ by a sequence of left-right decisions. For $i = 1, \ldots, \log_2 m - 1$ define the binary random variable Y_i that is 1 if at the i-th node the path P_ℓ continues left, and is 0 otherwise. Further, we use W_i^L and W_i^R to denote the

random variable that describes the total weight assigned to variables in X_{e_L} and X_{e_R}, respectively, where e_L and e_R, are the child-edges of the i-th node on path P_ℓ. Then,

$$E\left[\sum_{e\in P_\ell} w(X_e)\right] = E\left[\sum_{i=1}^{\log_2 m-1} Y_i\cdot W_i^L + (1-Y_i)\cdot W_i^R\right] = \sum_{i=1}^{\log_2 m-1} E[Y_i]\cdot E[W_i^L] + E[1-Y_i]\cdot E[W_i^R]$$
$$= \sum_{i=1}^{\log_2 m-1} \frac{1}{2}E[W_i^L + W_i^R] \geq \frac{1}{4}(\log_2 m - 1)\log_2 d \ .$$

Here, the second equality follows as Y_i is independent from W_i^L and W_i^R; the last inequality holds due to Lemma A.1.

B Omitted Proofs of Section 2

Proof of Lemma 2.1: Let (α, x) be an optimal solution of $\mathbb{SF}_\mathbb{IP}$ for a given graph G and a sequence of demands σ. We define subgraph F as the union of all edges of G whose value in x is equal to 1. Since constraints of type (1) ensure that the demands are satisfied, every pair (s_i, t_i) is connected in F. Without loss of generality we assume F is a forest, since otherwise removing any edge from a cycle of F provides a better solution for $\mathbb{SF}_\mathbb{IP}$ which contradicts the optimality of (α, x). Now, according to Lemma 2.3, there exists a connective list of subgraphs Q for $\langle G, F, \sigma\rangle$ such that every edge of F appears in at most $3\log^2 n$ subgraphs of Q where n is the size of the graph. We construct a solution (α', x') to $\mathbb{PC}_\mathbb{IP}$ in the following way:

Since Q is a connective list of subgraphs for $\langle G, F, \sigma\rangle$, every demand (s_i, t_i) is connected in $\cup_{j=1}^{i} Q_j$. Therefore, for every i, $Q_i \in \mathcal{H}_i$ holds. We set $\alpha' = (3\log^2 n)\alpha$, $x'^i_{Q_i} = 1$ for every $1 \leq i \leq k$, and $x'^i_H = 0$ for all other variables. Note that since $Q_i \in \mathcal{H}_i$ for all i, then constraints of type (6) are trivially satisfied. We define $\hat{x}_e$ for an edge $e \in E(G)$ as $\hat{x}_e = \sum_{i=1}^{k}\sum_{H\ni e} x'^i_H$. Since every edge of F appears in at most $3\log^2 n$ subgraphs of Q and x is a feasible solution for $\mathbb{SF}_\mathbb{IP}$, for every vertex v in $V(G)$ we have

$$\begin{aligned}
\frac{1}{b_v}\sum_{i=1}^{k}\sum_{H\subseteq G} \mathsf{deg}_H(v)x'^i_H &= \frac{1}{b_v}\sum_{e\in\delta(v)}\hat{x}_e \\
&\leq (3\log^2 n)\frac{1}{b_v}\sum_{e\in\delta(v)} x_e \\
&\leq (3\log^2 n)\alpha \\
&= \alpha'
\end{aligned}$$

and thus all constraints of type (4) are satisfied by x'. Moreover,

$$\begin{aligned}
\frac{1}{w(opt)}\sum_{i=1}^{k}\sum_{H\subseteq G} w(H)x'^i_H &= \frac{1}{w(opt)}\sum_{e\in E(G)} w(e)\hat{x}_e \\
&\leq (3\log^2 n)\frac{1}{w(opt)}\sum_{e\in E(G)} w(e) \\
&\leq (3\log^2 n)\alpha \\
&= \alpha'
\end{aligned}$$

and hence x' meets Constraint (5). Thus, x' is a feasible solution for $\mathbb{PC}_\mathbb{IP}$ and the proof is complete. ■

C Omitted Proofs of Section 3

Proof of Theorem 3.3: By Lemma 3.2, if the optimal solution of the IP is 1, then there exists a deterministic algorithm with competitive ratio $O(k \log m)$. We show one can use a doubling technique to obtain an $O(k \log m)$-competitive algorithm without such assumption.

We start by guessing $\alpha = 1$. We use Algorithm 1 to find an online solution. Since Algorithm 1 is $O(k \log m)$-competitive if $\alpha = 1$, there exists a constant β such that no packing constraint is violated by more than $\beta(k \log m)$. At each online stage j if updating the solution violates a packing constraint by more than $\beta(k \log m)$, then $\alpha > 1$. Thus we say we go to the next phase and do the following. First we divide every coefficient in packing constraints by 2. Note that this is equivalent to doubling α. Then we remove all previous solutions of the algorithm (we basically ignore them). We use Algorithm 1 again on the new IP from the beginning. We claim that we lose a factor of at most 4 due to these updates.

At each phase every packing constraint is violated by at most $\beta k \log m$. The total violation of each packing constraint after l phases is

$$\sum_{i=1}^{l} \beta 2^{i-1} k \log m \leq \beta 2^{i} k \log m \ . \tag{19}$$

Moreover our guess about α is not exact, since we know $2^{i-1} < \alpha \leq 2^i$. Thus we might violate each packing constraint by an additional factor of 2. Thus we can use Algorithm 1 to obtain a deterministic integral algorithm with competitive ratio $4\beta(k \log m)$. ■

D Doubling Method

In this section we discuss how we can address the issue that we might not know w_{opt} in advance. In particular by Lemma 3.2 if we are given the degree bounds and w_{opt}, the optimal solution for $\mathbb{PC_IP}$ is 1. Thus Algorithm 1 provides an online solution for $\mathbb{PC_IP}$ of $O(\log m)$, where m is the number of packing constraints. We show we can start with an initial value for w_{opt} and update it through several phases. We use Algorithm 1 at each phase. Eventually we lose a factor because of the number of phases we run Algorithm 1 and a factor of our approximation for w_{opt}.

More precisely, let w'_{opt} be our guess for w_{opt}. Initially we set $w'_{\text{opt}} = 1$. We write $\mathbb{PC_IP}$ according to w'_{opt} and use Algorithm $\mathbb{PC_IP}$ to obtain online solution $\mathbf{x}$. By Lemma 3.2 if $w'_{\text{opt}} \leq w_{\text{opt}}$, there exists a constant c such that every packing constraint is violated by a factor of at most $c \log m$ in solution $\mathbf{x}$. At any online step, if a packing constraint is going to be violated by more that $c \log m$, we stop the algorithm. Now we know $w'_{\text{opt}} > w_{\text{opt}}$, thus we update $w'_{\text{opt}} = w'_{\text{opt}} \times r$. We ignore the current solution $\mathbf{x}$. We write $\mathbb{PC_IP}$ by the updated value of w'_{opt} and do the same until no packing constraint is violated by more than $c \log m$.

Now we analyze the competitive ratio of the current solution. Let l denote the number of phases we updated w'_{opt} and run Algorithm 1.

Lemma D.1 *If $r \geq 2$ each degree constraint is violated by $O(l \log m)$, and the weight constraint is violated by $O(r \log m)$, where m is the number of packing constraints.*

Proof. Since at each phase no packing constraint is violated by more than $c \log m$, the total violation for every packing constraint is no more than $lc \log m$. Thus the total violation in every degree bound is $O(l \log m)$.

At each phase, the packing constraint corresponding to weights is violated by $c \log m$, however w'_{opt} is changing as well. The total violation of the weights is $\sum_{i=0}^{l-1} r^i c \log m$, since at each phase i, $w'_{\text{opt}} = r^{i-1}$.

Since $r \geq 2$, $\sum_{i=0}^{l-1} r^i c \log m < 2r^{l-1} = 2w'_{\text{opt}}$. Therefore the packing constraint corresponding to weights is also violated by no more than $2c \log m$. However $\frac{1}{r} w'_{\text{opt}} < w_{\text{opt}} \leq w'_{\text{opt}}$. Thus the total weight is no more than $2rc \log m w_{\text{opt}}$. □

The following two lemmas follow directly from Lemma D.1.

Lemma D.2 *Given graph G and degree bounds. let w_{opt} denote the optimal weight for* ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST. *There exists an online algorithm for $\mathbb{PC}_\mathbb{IP}$ that violates the weight packing constraint by $O(\log n)$ and violates each degree bound packing constraints by $O(\log n \log w_{opt})$.*

Proof. Let $r = 2$. The number of phases $l = O(\log w_{\text{opt}})$. Thus by Lemma D.1 there exists an online algorithm for $\mathbb{PC}_\mathbb{IP}$ that violates the weight packing constraint by $O(\log n)$ and violates each degree bound packing constraints by $O(\log n \log w_{\text{opt}})$. □

Lemma D.3 *Given graph G and degree bounds. let w_{opt} denote the optimal weight for* ONLINE EDGE-WEIGHTED DEGREE-BOUNDED STEINER FOREST. *There exists an online algorithm for $\mathbb{PC}_\mathbb{IP}$ that violates the weight packing constraint and each degree bound packing constraints by $O(\log n \frac{\log w_{opt}}{\log \log w_{opt}})$.*

Proof. Let $r = \frac{\log w_{\text{opt}}}{\log \log w_{\text{opt}}}$. The number of phases $l = O(\frac{\log w_{\text{opt}}}{\log \log w_{\text{opt}}})$. Thus by Lemma D.1 there exists an online algorithm for $\mathbb{PC}_\mathbb{IP}$ that violates the weight packing constraint and each degree bound packing constraints by $O(\log n \frac{\log w_{\text{opt}}}{\log \log w_{\text{opt}}})$. □

E A Polynomial Algorithm for Finding the Path with the Minimum Cost

Here we present a polynomial time method to find a set S that satisfies $j+1$-th constraint and minimizes $\tau_S(\mathcal{S}(j))$. In our problem we can assume every two endpoints of previous demands have been contracted into one node. With this assumption the set S is in form of a path from s_j to t_j. The method we use is dynamic programming. Let dp be a dynamic table such that $dp(l, v, w)$ denotes the minimum degree cost of a path with length $l < n$ and weight w from s_j to vertex $v \in V$. The degree cost of a path P is the sum of the terms corresponding to degree constraints in the calculation of $\tau_P(\mathcal{S}(j))$. The set S we are looking for has some length l' and some weight w'. Knowing these two parameters, we can calculate $\tau_S(\mathcal{S}(j))$ using w' and $dp(l', t_j, w')$.

If the maximum weight of an edge be a polynomial of n then the number of entries in the dynamic table dp is polynomial. Now we briefly explain how to calculate the value of each $dp(l, v, w)$. Initially set $dp(0, s_j, 0)$ to zero and $dp(l, v, w)$ to infinity for other entries. At every moment each entry with cost less than infinity represents a path from s_j. At every iteration we update new entries by adding one vertex to the existing paths.

We can now find the path minimizing $\tau_S(\mathcal{S}(j))$ by looking over $dp(l', t_j, w')$, for all $0 \leq l' \leq n$ and $0 \leq w' \leq n \max_{e \in E}\{w(e)\}$. Since the size of dp and the number of iterations are polynomials of n, the whole method runs in polynomial time.

F Rounding Lemma

For a given tree $T = (E(T), V(T))$, let $\pi = \langle \pi_1, \ldots, \pi_n \rangle$ be a permutation on $V(T)$, and $\mathcal{F}$ be a collection of subsets of $E(T)$. Suppose a probability p_{π_i, π_j} is assigned to every $\mathcal{F}_{\pi_i, \pi_j}$ in a way that $\sum_{j=1}^{i-1} p_{\pi_i, \pi_j} = 1$,

for every $1 < i \leq n$. We define the load of p on an edge e as $L_p(e) := \sum_{(i,j):e \in \mathcal{F}_{\pi_i,\pi_j}} p_{\pi_i,\pi_j}$ and the overall load L_p as $L_p := \max_{e \in E(t)}\{L_p(e)\}$.

The following lemma states the existence of a rounding q for p, which with a high probability has a bounded load with respect to L_p and $\log n$.

Lemma F.1 *There exists a new assignment $q_{\pi_i,\pi_j} \in \{0,1\}$ to every $\mathcal{F}_{\pi_i,\pi_j}$, such that $\sum_{j=1}^{i-1} q_{\pi_i,\pi_j} = 1$, for every $1 < i \leq n$, and the new load L_q is at most $O(\max\{L_p, \log n\})$ with a high probability.*

Proof. For every $1 < i \leq n$ we need to set at least one of $q_{\pi_i,\pi_1}, \ldots, q_{\pi_i,\pi_{i-1}}$ to 1. To do so, we select one of the q_{π_i,π_j}'s using the following random process. Take a random number $r \in [0,1]$. Let j be the smallest index for which $\sum_{k=1}^{j} p_{\pi_i,\pi_k} \geq r$. Set q_{π_i,π_j} to 1 and the remaining to 0. Note that the probability of every q_{π_i,π_j} being 1 is exactly p_{π_i,π_j}.

Using this random process, the expected load of q on every edge e remains the same as the load of p on that edge, because

$$\begin{aligned} E[L_q(e)] &= E\Big[\sum_{(i,j):e \in \mathcal{F}_{\pi_i,\pi_j}} q_{\pi_i,\pi_j}\Big] \\ &= \sum_{(i,j):e \in \mathcal{F}_{\pi_i,\pi_j}} E[q_{\pi_i,\pi_j}] \\ &= \sum_{(i,j):e \in \mathcal{F}_{\pi_i,\pi_j}} Pr[q_{\pi_i,\pi_j} = 1] \\ &= \sum_{(i,j):e \in \mathcal{F}_{\pi_i,\pi_j}} p_{\pi_i,\pi_j} = L_p(e) \ . \end{aligned}$$

Moreover, $L_q(e)$ is in fact the summation of a number of binary random variables which are not positively correlated [3]. Therefore, this summation can be upper bounded by the Chernoff bound:

$$Pr[L_q(e) > (1+\delta)L_p(e)] \leq e^{-\frac{\delta L_p(e)}{3}} \ .$$

In the above variation of the Chernoff bound, $\delta > 1$. Mention that to achieve a small enough probability [4], it suffices for δ to be at least $\Omega(\log n / L_p(e))$. Finally, we use the Union bound to show that $L_q \in O(\max\{L_p, \log n\})$ with a high probability.

$$\begin{aligned} Pr[L_q \in O(\max\{L_p, \log n\})] &= Pr[\forall e \in E(T) : L_q(e) \in O(\max\{L_p(e), \log n\})] \\ &= 1 - Pr[\exists e \in E(T) : L_q(e) \notin O(\max\{L_p(e), \log n\})] \\ &\geq 1 - \sum_{e \in E(T)} Pr[L_q(e) \notin O(\max\{L_p(e), \log n\})] \\ &\geq 1 - \sum_{e \in E(T)} Pr[L_q(e) > c\big(1 + \frac{\log n}{L_p(e)}\big)L_p(e)] \\ &\geq 1 - \sum_{e \in E(T)} O(\frac{1}{n^2}) = 1 - O(\frac{1}{n}) \ . \end{aligned}$$

□

[3] In particular, $q_{\pi_{i_1},\pi_{j_1}}$ and $q_{\pi_{i_2},\pi_{j_2}}$ are independent for $i_1 \neq i_2$, and are negatively correlated for $i_1 = i_2$ and $j_1 \neq j_2$.

[4] At most $O(\frac{1}{n^2})$.

G Reduction from weight guarantee to edge-wise guarantee

Lemma G.1 *Let $\mathcal{R}$ be a subspace of $\mathbb{R}^n$ that contains all points of $\mathbb{R}^n$ with non-negative coordinates and P be a convex set of points in $\mathcal{R}$. If for every point $\hat{x} = (x_1, x_2, \ldots, x_n) \in \mathcal{R}$ there exists a point $\hat{p} \in P$ such that*

$$\hat{p}.\hat{x} \leq k \sum_{i=1}^{n} x_i$$

then P contains a point $\hat{r}$ such that $\max_{i=1}^{n} r_i \leq k$.

Proof. We define P' as the set of all points in $\mathbb{R}^n$ whose all indices are greater than or equal to the corresponding indices of a point in P. In other words

$$P' = \{\hat{p} \in \mathbb{R}^n | \exists \hat{q} \in P \text{ such that } p_i \geq q_i \text{ for all } 1 \leq i \leq n\}.$$

We show in the rest that $(k, k, \ldots, k) \in P'$ which immediately implies the lemma. To this end, suppose for the sake of contradiction that $(k, k, \ldots, k) \notin P'$. Note that, since P is a convex set, so is P'. Therefore, there exists a hyperplane that separates all points of P' from point $(k, k, \ldots, k)$. More precisely, there exist coefficients $h_0, h_1, \ldots, h_n$ such that

$$\sum_{i=1}^{n} h_i p_i > h_0 \tag{20}$$

for all points $\hat{p} \in P'$ and

$$\sum_{i=1}^{n} h_i k < h_0. \tag{21}$$

Due to the construction of P' we are guaranteed that all coefficients $h_1, h_2, \ldots, h_n$ are non-negative numbers since otherwise for any index i such that $h_i < 0$ there exists a point in P' whose i'th index is large enough to violate Inequality (20). Now let $\hat{x} = (h_1, h_2, \ldots, h_n)$. By Inequalities (20) and (21) we have

$$\hat{p}.\hat{x} = \sum_{i=1}^{n} x_i p_i = \sum_{i=1}^{n} h_i p_i > h_0 \geq k \sum_{i=1}^{n} h_i = k \sum_{i=1}^{n} x_i$$

for every $\hat{p} \in P'$ which means there is no $\hat{p} \in P$ such that $\hat{p}.\hat{x} \leq k \sum_{i=1}^{n} x_i$. This contradicts the assumption of the lemma. □